

## CONCURRENT SCHEDULES: SHORT- AND LONG-TERM EFFECTS OF REINFORCERS

JASON LANDON, MICHAEL DAVISON, AND DOUGLAS ELLIFFE

UNIVERSITY OF AUCKLAND, NEW ZEALAND

Five pigeons were trained on concurrent variable-interval schedules in a switching-key procedure. The overall rate of reinforcement was constant in all conditions, and the ratios of reinforcers obtainable on the two alternatives were varied over seven levels. Each condition remained in effect for 65 sessions, and the last 50 sessions of data from each condition were analyzed. The most recently obtained reinforcer had the largest effect on current preference, but each of the eight previously obtained reinforcers had a small measurable effect. These effects were larger when the reinforcer ratio was more extreme. A longer term effect of reinforcement was also evident, which changed as a function of the reinforcer ratio arranged. More local analyses showed regularities at a reinforcer-by-reinforcer level and large transient movements in preference toward the just-reinforced alternative immediately following reinforcers, followed by a return to stable levels that were related to the reinforcer ratio in effect. The present data suggest that the variables that control choice have both short- and long-term effects and that the short-term effects increased when the reinforcer ratios arranged were more extreme.

*Key words:* concurrent schedules, choice, generalized matching, reinforcement, key peck, pigeon

Most previous research on choice has focused on the description and understanding of steady-state performance. A common approach has been to use concurrent variable-interval (VI) VI schedules. An independent variable, often the reinforcer ratio, is held constant until behavior stabilizes. This typically requires about 15 to 30 sessions of training, with data from only the last few sessions being analyzed.

This research has shown that the ratio of responses or time allocated to the alternatives follows the ratio of reinforcers obtained from the alternatives. The generalized matching law (Baum, 1974) describes choice well in a variety of situations (see Davison & McCarthy, 1988). The logarithmic version of the generalized matching law, where the log of the ratio of responses emitted on the two alternatives is a linear function of the log of the ratio

of reinforcers obtained from the two alternatives, is given by

$$\log\left(\frac{B_1}{B_2}\right) = a \log\left(\frac{R_1}{R_2}\right) + \log c, \quad (1)$$

where  $B_1$ ,  $B_2$ ,  $R_1$ , and  $R_2$  are, respectively, the numbers of responses emitted on and reinforcers obtained from Alternatives 1 and 2. The parameter  $a$  is sensitivity to reinforcement (Lobb & Davison, 1975), and measures the change in response ratios resulting from a change in the reinforcer ratios.  $\log c$  is inherent bias, and measures any constant proportional preference for one alternative over the other.

A large body of research has shown that the value of  $a$  in Equation 1 is typically about 0.80 to 0.85 when response measures of behavior allocation are used, whereas when time-allocation measures are used this value is slightly higher, 0.90 to 0.95 (Baum, 1979; Davison & McCarthy, 1988; Taylor & Davison, 1983; Wearden & Burgess, 1982). Thus, undermatching is the standard result on concurrent VI VI schedules. The success of the generalized matching law in describing data from a large number of experiments seems to imply that reinforcers aggregated over relatively long periods of time control behavior.

More recently, investigators have begun to examine behavior when experimental conditions change more frequently than the typical 15 to 30 sessions. Davison and Hunter (1979)

The experiments reported here were completed in partial fulfillment of a doctoral degree at the University of Auckland by the first author. We thank William Baum for his continuing contribution to this research, Mick Sibley for looking after the subjects, and members of the Experimental Analysis of Behavior Research Unit for helping run the experiments. The data used in the analyses are too extensive to include in an appendix, and may be obtained from the authors on receipt of a 100-Mb zip disk.

Reprints may be obtained from any author, Department of Psychology, University of Auckland, Private Bag 92019, Auckland, New Zealand (e-mail: j.landon@auckland.ac.nz, m.davison@auckland.ac.nz, d.elliffe@auckland.ac.nz).

investigated concurrent VI VI schedule performance when the arranged reinforcer ratio was changed every six sessions. They started with six sessions of concurrent extinction VI 60 s, and the reinforcer ratio was progressively increased, ending with six sessions of concurrent VI 60 s extinction. This sequence was reversed and repeated twice, giving two ascending and two descending sequences. Davison and Hunter used generalized matching analyses to assess the control exerted over current behavior by the previous sessions' reinforcer ratios. The first, third, and sixth sessions following a reinforcer-ratio change were examined. They showed that the previous reinforcer ratio had a large effect on preference in the first session following a change. There was still a measurable effect three sessions after a change, but this effect was no longer detectable six sessions after a change. Sensitivity to the current reinforcer ratio reached close to its asymptotic level (about 0.80) after about six sessions following a change in ratio.

Hunter and Davison (1985) examined behavior in a more rapidly changing procedure. They exposed pigeons to either concurrent VI 60-s VI 240-s or concurrent VI 240-s VI 60-s schedules. The schedule that was in effect changed daily according to a 31-step pseudorandom binary sequence. Thus, which reinforcer ratio was to be in effect in a given session was not predictable. Like Davison and Hunter (1979), Hunter and Davison found that reinforcer ratios obtained in only the previous two or three sessions affected current performance.

Schofield and Davison (1997) systematically replicated and extended Hunter and Davison's (1985) research. However, Schofield and Davison found that, with repeated exposures to pseudorandom binary sequences, measurable effects of previous sessions' reinforcer ratios were no longer discernible, even in the immediately following session. Their result suggests that under extended exposure to frequently changing reinforcer ratios, the control of behavior had become much more localized.

Davison and Baum (2000) adapted a procedure introduced by Belke and Heyman (1994) to study performance in even more frequently changing environments. Seven different reinforcer ratios (27:1, 9:1, 3:1, 1:1, 1:3,

1:9, and 1:27) were presented to the subjects in a random order in each experimental session. Each of these components was separated by a 10-s blackout of the keylights, and across conditions the number of reinforcers per component was varied. They found moderate sensitivity ( $a = 0.60$ ) to the reinforcer ratio in the current component after as few as six to eight reinforcer deliveries in that component. That is, control by the current component developed very rapidly and became very localized. There was, however, no effect on behavior of component length.

Considered with steady-state results, research that has examined behavior in more frequently changing conditions (Davison & Baum, 2000; Davison & Hunter, 1979; Hunter & Davison, 1985; Schofield & Davison, 1997) indicated that the speed of adjustment of choice to changing reinforcer contingencies is itself a function of how quickly or how often those contingencies change, and of the amount of previous exposure to those changes. That is, sensitivity to previous reinforcer ratios declined more quickly, and sensitivity to the current reinforcer ratio increased more quickly, when contingencies changed more quickly. Indeed, in Davison and Baum's procedure, in which reinforcer ratios changed several times within each session, individual reinforcers produced strikingly regular and predictable changes in preference, each reinforcer reliably moving preference toward the alternative from which it was obtained. Successive reinforcers obtained from the same alternative had diminishing effects on behavior. In contrast, irrespective of its sequential position, a single reinforcer obtained from the opposite alternative following a sequence of reinforcers from the same alternative (termed a *disconfirmation* by Davison & Baum, and here) had a much larger effect on behavior. Davison and Baum suggested that their results could be modeled using a single leaky accumulator, allowing reinforcer deliveries to have immediate effects on behavior that lasted for a matter of minutes—local control by reinforcers. This was generally confirmed by Davison and Baum (2002).

Simple frequency of change in reinforcing environments, however, may not be the only variable that controls the speed with which behavior changes. Landon and Davison

(2001) replicated and extended Davison and Baum's (2000) results using the same basic procedure, but they varied the range over which component reinforcer ratios changed during a session. They showed that a given reinforcer had a larger effect when the range of reinforcer ratios arranged was large (e.g., 27:1 through 1:27) than when it was small (e.g., 1.5:1 through 1:1.5). Nevertheless, the same basic effects, including diminishing effects of successive reinforcers on the same alternative and large effects of disconfirmations, were seen.

Two possible implications of the local effects of individual reinforcers reported by Davison and Baum (2000) and Landon and Davison (2001) are immediately apparent. The first, suggested by our outline of the literature, is that they represent an extreme form of the increasing rate of behavioral change produced by increasingly rapid changes in contingencies. Thus, local effects of reinforcers may not occur in stable-state concurrent VI VI performance, and behavioral control in this procedure may result from long-term aggregates of obtained reinforcers. The second is that similar effects are indeed present in conventional, steady-state concurrent VI VI performance, but have been overlooked because the relevant detailed data have not been collected and analyzed. If the latter is true, then it is possible that steady-state performance may result only from such local effects, or that it may arise from a combination of local and more long-term reinforcer effects.

To address this question, we conducted a conventional steady-state switching-key concurrent VI VI experiment. In each condition, we arranged one of the seven reinforcer ratios used by Davison and Baum (2000). To ensure that sufficient data were collected to allow analyses at a local level, each condition was in effect for 65 sessions, and the last 50 sessions of data were analyzed.

## METHOD

### *Subjects*

The subjects were 5 homing pigeons numbered 131, 132, 134, 135, and 136. Another pigeon, numbered 133, died during the present experiment, and no data obtained from

that subject are included here. The subjects were maintained at  $85\% \pm 15$  g of their free-feeding body weights by postsession feeding of appropriate amounts of mixed grain. Water and grit were freely available at all times.

### *Apparatus*

Each pigeon was housed in a cage 380 mm high, 380 mm wide, and 380 mm deep. The back, left, and right walls of each cage were constructed of sheet metal; the top, floor, and front wall consisted of metal bars. Each cage contained two wooden perches, the first mounted 95 mm from and parallel to the front wall, and the second mounted 95 mm from and parallel to the right wall.

The right wall of each cage contained three translucent response keys, 20 mm in diameter, centered 100 mm apart and 200 mm above the perches. The center key remained dark and inoperative throughout. The left key could be lit yellow, and the right key could be lit either red or green. Both keys, when lit, could be operated by pecks exceeding a force of approximately 0.1 N. A hopper containing wheat was located behind an aperture (50 mm by 50 mm) situated 145 mm below the center key. During reinforcer delivery, the hopper was raised to the aperture and illuminated for 4 s and the keylights were extinguished. All experimental events were arranged on an IBM® PC-compatible computer running MED-PC® software, located in a room remote from the experimental cages. The computer recorded the time, at 10-ms resolution, at which every event occurred in experimental sessions.

### *Procedure*

A switching-key (Findley, 1958) concurrent-schedule procedure was used. Sessions began with the left (switching) key lit yellow, and the right (main) key lit either red or green with equal probability. Reinforcers were scheduled according to a single exponential VI 30-s schedule ( $p = .033$  per second). Once a reinforcer had been arranged, it was allocated to either the red or green alternative according to the probability for each condition, as shown in Table 1. Reinforcers were dependently scheduled (Stubbs & Pliskoff, 1969), meaning that once a reinforcer had been arranged for one alternative, no further reinforcers were arranged until that reinforcer

Table 1

Sequence of experimental conditions, relative reinforcer probability (shown as probability of reinforcement on the red alternative), and the reinforcer ratio for each of the eight conditions. The overall probability of reinforcement per second was constant at .033 throughout.

Condition	Relative reinforcer probability $p(R)$	Reinforcer ratio (R : G)
1	.5000	1:1
2	.9643	27:1
3	.2500	1:3
4	.9000	9:1
5	.1000	1:9
6	.7500	3:1
7	.0357	1:27
8	.2500	1:3

had been obtained. A 2-s changeover delay (Herrnstein, 1961) prevented responses from producing an arranged reinforcer until 2 s had elapsed since the last switching-key response.

The sequence of experimental conditions is shown in Table 1. Across conditions, the overall rate of reinforcement was constant, and the red:green reinforcer ratio was varied between 1:27 and 27:1 through seven logarithmically equally spaced levels. Condition 8 replicated Condition 3. No stability criterion was in effect. However, 65 sessions were conducted for each condition to ensure that sufficient data were collected to allow analysis of particular sequences of reinforcers. The data from the last 50 sessions of each condition were used in the analyses. Sessions were conducted daily, and ended in blackout after 80 reinforcers had been obtained or after 42 min had elapsed, whichever occurred first.

## RESULTS

The contribution of recent reinforcers to current preference was examined as follows. First, for consistency with Davison and Baum (2000) and Landon and Davison (2001), a moving temporal window of the eight most recently obtained reinforcers was used. There were therefore 256 distinct sequences of red and green reinforcers in that window. Beginning with the eighth reinforcer in a session, red and green response numbers after each successive reinforcer (up to the next reinforcer) were aggregated according to which of those 256 sequences in the last eight re-

inforcers they followed, and a log red:green response ratio was calculated as a measure of current preference.

Not every sequence occurred in every condition. For example, Condition 7 arranged a red:green reinforcer ratio of 1:27. A sequence of eight successive red-key reinforcers was therefore very unlikely, and in fact never occurred in the 50 sessions analyzed for that condition. Also, no responses were emitted on one or the other alternative after some sequences. Because of these factors, it was not possible to calculate a log response ratio for every reinforcer sequence in every condition.

For those sequences of reinforcers that did occur and for which a log response ratio could be calculated, the contribution of each of the preceding eight reinforcers to the current log response ratio was measured by means of the following general linear model:

$$\log\left(\frac{B_R}{B_G}\right) = \log k + \sum_{j=0}^7 \begin{cases} R_j = R: & +b_j \\ R_j = G: & -b_j \end{cases} \quad (2)$$

In Equation 2, the subscript  $j$  denotes reinforcer lags in the preceding sequence of eight reinforcers, so that  $R_0$  is the most recent reinforcer. The coefficients  $b_j$  are log (Base 10) response ratios, and represent the amount of current preference attributable to the reinforcer at lag  $j$ . If that reinforcer was obtained from the red alternative,  $b_j$  is added, because the log response ratio should move in a positive direction. Similarly,  $b_j$  is subtracted if the reinforcer at lag  $j$  was obtained from green. We call  $b_j$  the *log reinforcer effect*. It is conceptually, but not quantitatively, analogous to sensitivity to reinforcement ( $a$  in Equation 1) at each lag. The constant  $\log k$  is also a log response ratio; it measures the residual amount of current preference that is not attributable to any of the eight most recent reinforcers. This constant therefore combines the more distant effects of prior reinforcers with any inherent bias (i.e.,  $\log c$  in Equation 1).

We obtained the best fitting least squares estimates of  $b_j$  and  $\log k$  by fitting Equation 2 to the log response ratios following each eight-reinforcer sequence using Quattro Pro® v. 8's Optimizer function. This analysis was carried out separately for each condition and for each subject. Figure 1 shows the results of these analyses. The log reinforcer effect ( $b_j$ )

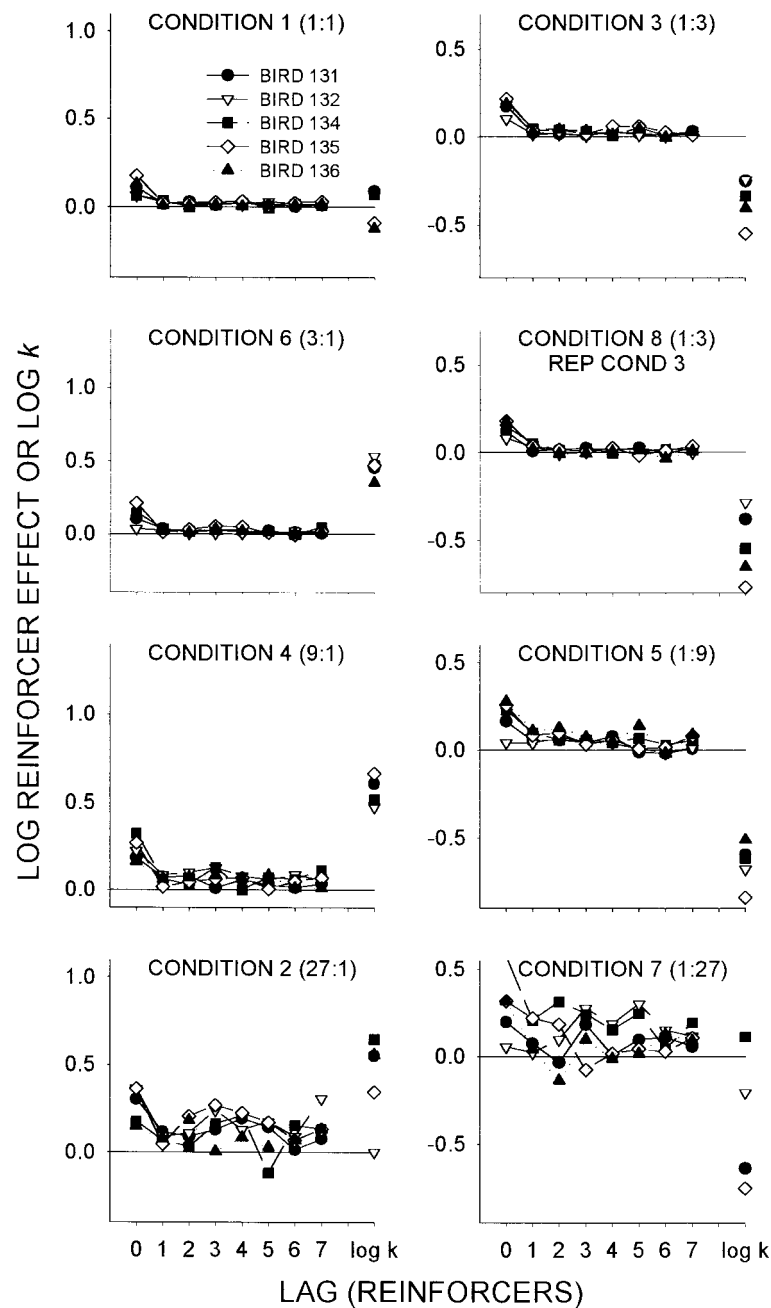


Fig. 1. Log reinforcer effect on current preference of each of the previous eight reinforcers for each subject plotted as a function of reinforcer lag (0 being the most recently obtained reinforcer). Also shown are the values of  $\log k$  (see Equation 2) for each subject.

of each of the preceding reinforcers is plotted as a function of reinforcer lag (Lag 0 is the most recent reinforcer) for each subject and each condition. The constant ( $\log k$ ) is also shown for each subject and condition.

Figure 1 shows four major effects. First, the most recently obtained reinforcer had by far the largest effect on current behavior. Second, in each condition, reinforcers beyond lag zero had similar effects on current behav-



ior. Third, reinforcers had larger effects on behavior in conditions in which the reinforcer ratio was more extreme. Fourth, the constant ( $\log k$ ) was more extreme when the reinforcer ratio was more extreme. Figure 1 shows that the results in the original Condition 3 and its replication (Condition 8) were similar. Therefore, the data from Condition 3 were used rather than those from the replication in this and subsequent analyses. A two-way repeated measures analysis of variance (ANOVA) was used to confirm the effects of reinforcer ratio and of the sequential position of the reinforcer on the log reinforcer effect. The ANOVA showed a significant effect of the reinforcer ratio,  $F(5, 168) = 21.69$ ,  $p < .05$ , but no significant effect of the sequential position of the reinforcer,  $F(6, 168) = 1.98$ ,  $p > .05$ .

To investigate further the effect of reinforcer ratio on log reinforcer effect, a Page planned comparisons<sup>1</sup> test (Marascuilo & McSweeney, 1977; Page, 1963) weighting mean ranks by quadratic orthogonal polynomials was applied to the mean data, across conditions, from the individual subjects. A significant quadratic effect was found ( $k = 7$  conditions,  $N = 8$  reinforcers;  $z = -6.31$ ,  $p < .05$ ), showing that the effect of a reinforcer increased as the arranged reinforcer ratio became more extreme. The upper panel in Figure 2 plots these mean data as a function of both the arranged log reinforcer ratio and reinforcer lag, and the lower panel plots log reinforcer effect at Lag 0 and the mean log reinforcer effect from Lags 1 through 7 as a function of the arranged log reinforcer ratio. These graphs show the relation between reinforcer effect and reinforcer ratio, and the upper panel shows that reinforcers beyond Lag 0 did not have any differential effects on current behavior.

As mentioned earlier, the constant ( $\log k$ ) in Equation 2 measures the residual current preference once the effects of each of the eight previous reinforcers are removed, and so might be best viewed as the effects of more distant aggregations of reinforcers in the current condition. A one-way ANOVA was used

to confirm that  $\log k$  did in fact change as the reinforcer ratio changed,  $F(6, 28) = 26.87$ ,  $p < .05$ .

The use of an eight-reinforcer window in the above analyses shows how previous reinforcers affect current preference. Given the large effects of reinforcers at Lag 0 and the smaller effects of reinforcers beyond Lag 0, however, it is possible that a smaller window might suffice to describe the data. Therefore, we repeated these analyses and varied the window size from one to eight reinforcers. Values of log reinforcer effect were not changed in any systematic way by this variation. Figure 3 shows the values of  $\log k$  obtained from these analyses plotted as a function of the window size for each subject in each condition. Nonparametric tests for trend (Ferguson, 1966) confirmed that, in all conditions except Condition 1 (1:1 reinforcer ratio),  $\log k$  became significantly less extreme as the window size was increased. The  $z$  scores obtained for the trend tests in each condition are shown in Figure 3. Decreases in  $\log k$  were greater in conditions in which the reinforcer ratio was more extreme. The values of  $\log k$  support the earlier analysis suggesting that previous reinforcers did have an effect on current preference, and that this effect was greater when the reinforcer ratio was more extreme. Moreover, the trends suggest that control may have been more local in conditions in which the reinforcer ratio was more extreme.

A local analysis similar to those reported by both Davison and Baum (2000) and Landon and Davison (2001) broke the data into log response ratios emitted in interreinforcer intervals following every sequence of reinforcers obtained in a condition. A sliding window of nine reinforcers was used. Thus, before the first reinforcer in a sequence, one log response ratio could be calculated. After the first reinforcer and before the second, two log response ratios were available (one following a red reinforcer, and one following a green reinforcer). After two reinforcers in a sequence, four log response ratios were available, one for each possible two-reinforcer sequence, and so on.

Figure 4 shows the log response ratio emitted following sequences of red or green reinforcers obtained in succession, and the effects of a single disconfirmation at each

<sup>1</sup> The use of a planned comparisons test is justified by Landon and Davison's (2001) results, which allow an a priori prediction that log reinforcer effect will be greater when the reinforcer ratio is more extreme.




Fig. 2. The upper panel shows log reinforcer effect plotted as a function of reinforcer lag and the log arranged reinforcer ratio, averaged over the 5 subjects. The lower panel shows log reinforcer effect at Lag 0 and the mean value across Lags 1 through 7, both as a function of the log arranged reinforcer ratio.

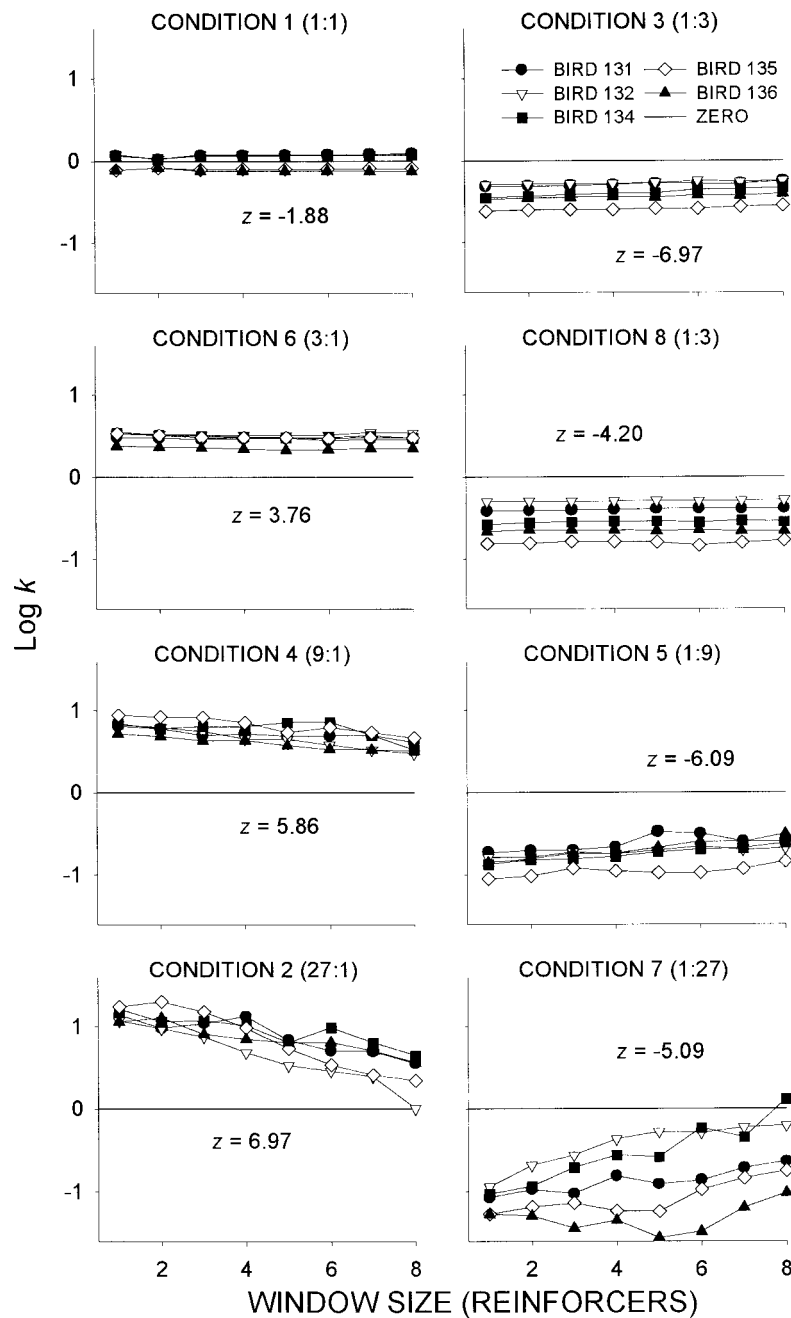


Fig. 3.  $\log k$  in each condition for each subject plotted as a function of the size of the moving window used in the analyses. Also shown are the  $z$  scores obtained from nonparametric tests for trend in each condition.

sequential position in each condition. Data were omitted from the graphs if fewer than 30 responses were recorded at either alternative. Local effects, similar to those reported by Davison and Baum (2000) and Landon

and Davison (2001), were evident in the present data. Specifically, successive reinforcers obtained from the same alternative generally moved preference towards the alternative from which they were obtained, irrespective



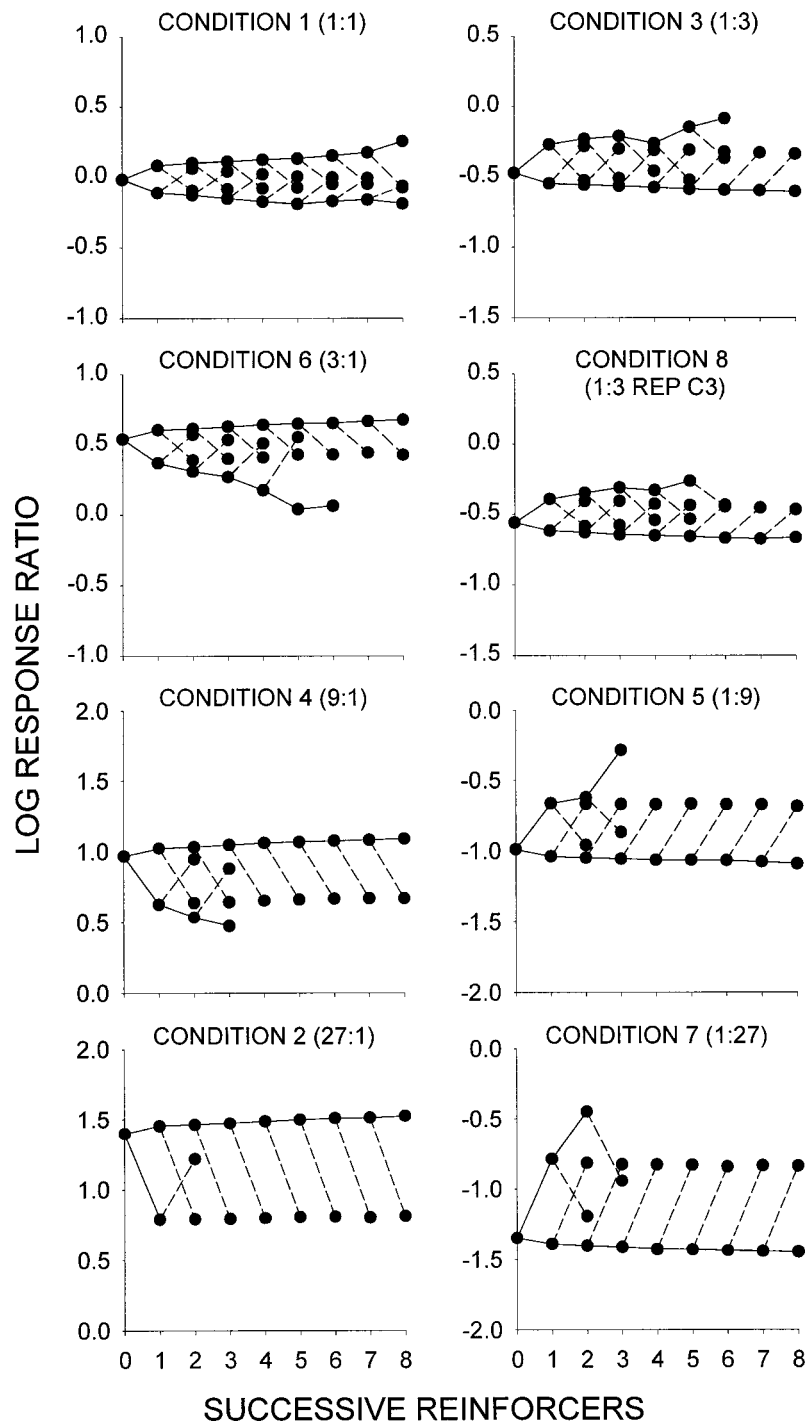


Fig. 4. Log response ratios emitted in interreinforcer intervals following successive same-alternative reinforcers (solid lines). The broken lines join disconfirmations, where a reinforcer was obtained from the other alternative following sequences of successive same-alternative reinforcers. A sliding window of nine reinforcers was used throughout, and data were omitted if fewer than 30 responses were recorded at either alternative.

of the reinforcer ratio arranged in that condition. Moreover, disconfirmations had comparatively large effects on preference.

One point of difference between the present results and those reported earlier by Davison and Baum (2000) and Landon and Davison (2001) is the asymmetry in the tree structures in Figure 4. Successive same-alternative reinforcers obtained at the alternative providing the lower rate of reinforcement had larger effects on the log response ratio than those obtained at the alternative providing the higher rate of reinforcement. In each condition, there was also an apparent limit on how far preference moved following sequences of same-alternative reinforcers obtained at the alternative providing the higher rate of reinforcement, and these limits were ordered in the same way as the reinforcer ratios. Moreover, the tree structures as a whole were biased toward the alternative that provided the higher rate of reinforcement.

The results presented in Figure 4 show behavior emitted in successive interreinforcer intervals. An even more local analysis was used to examine behavior during interreinforcer intervals in each condition. To do this, the data were collated across subjects into 5-s time bins following all possible two-reinforcer sequences. Separate log response ratios were calculated for each 5-s bin following each of the four sequences of reinforcers; Figure 5 shows the log red:green response ratio in each 5-s bin for each sequence plotted as a function of time since reinforcement.

The missing leftmost points in Figure 5 show that in many cases (18 of 32), preference in the first 5-s bin following a reinforcer was exclusive to the alternative from which a reinforcer had just been obtained. A number of consistencies are evident across conditions. First, a single reinforcer in each condition resulted in a large transient shift in preference towards the alternative at which it was obtained. For example, in Condition 1 (1:1 reinforcer ratio), the effects of a red- or green-alternative reinforcer were symmetrical: In the first two bins after reinforcement, a strong preference was evident for the just-reinforced alternative, and thereafter the log response ratio quickly returned to a stable level close to zero. In other conditions similar transient movements in preference were also evident. Across conditions, however, prefer-

ence stabilized at levels that changed as a function of the log reinforcer ratio arranged in that condition, and thus might be considered analogous to log  $k$  above.

Some representation of the quadratic effect of the reinforcer ratio on log reinforcer effect shown in Figure 2 might also be expected in Figure 5. Indeed, a similar effect was evident in the duration of the transient movement in preference. If Condition 1 (1:1) is considered, preference had, in all cases, returned to its stable level by the third or fourth 5-s bin following a reinforcer delivery. In Condition 4 (9:1), preference did not return to its stable level until the sixth or seventh 5-s bin following a reinforcer delivery on the more frequently reinforced alternative, and in Condition 2 (27:1) this was not the case until about the ninth 5-s bin. Across conditions, the duration of the transient movement in preference increased as the reinforcer ratio increased.

Figure 6 shows the mean number of responses emitted in each successive 5-s bin following each of the four sequences of reinforcers as a function of time since reinforcement. Data were omitted when a particular time bin was reached fewer than five times. Response rates were similar across conditions, and it is clear that they were unaffected by whether the most recent reinforcer was from the same or the opposite alternative to the preceding reinforcer. In conditions in which reinforcement was unequal, however, response rates differed in the first two bins after a reinforcer delivery. In all cases, the response rate following a reinforcer obtained from the leaner alternative was substantially lower than that following a reinforcer at the richer alternative. These response rates were about equal by the third 5-s bin after a reinforcer delivery.

## DISCUSSION

The present results show that the local effects of reinforcement evident when contingencies changed very rapidly (approximately every 10 reinforcers; Davison & Baum, 2000; Landon & Davison, 2001) were also present in a steady-state procedure (when contingencies were changed every 65 sessions). Thus, the regular effects on preference of individual reinforcers seen in the previous research

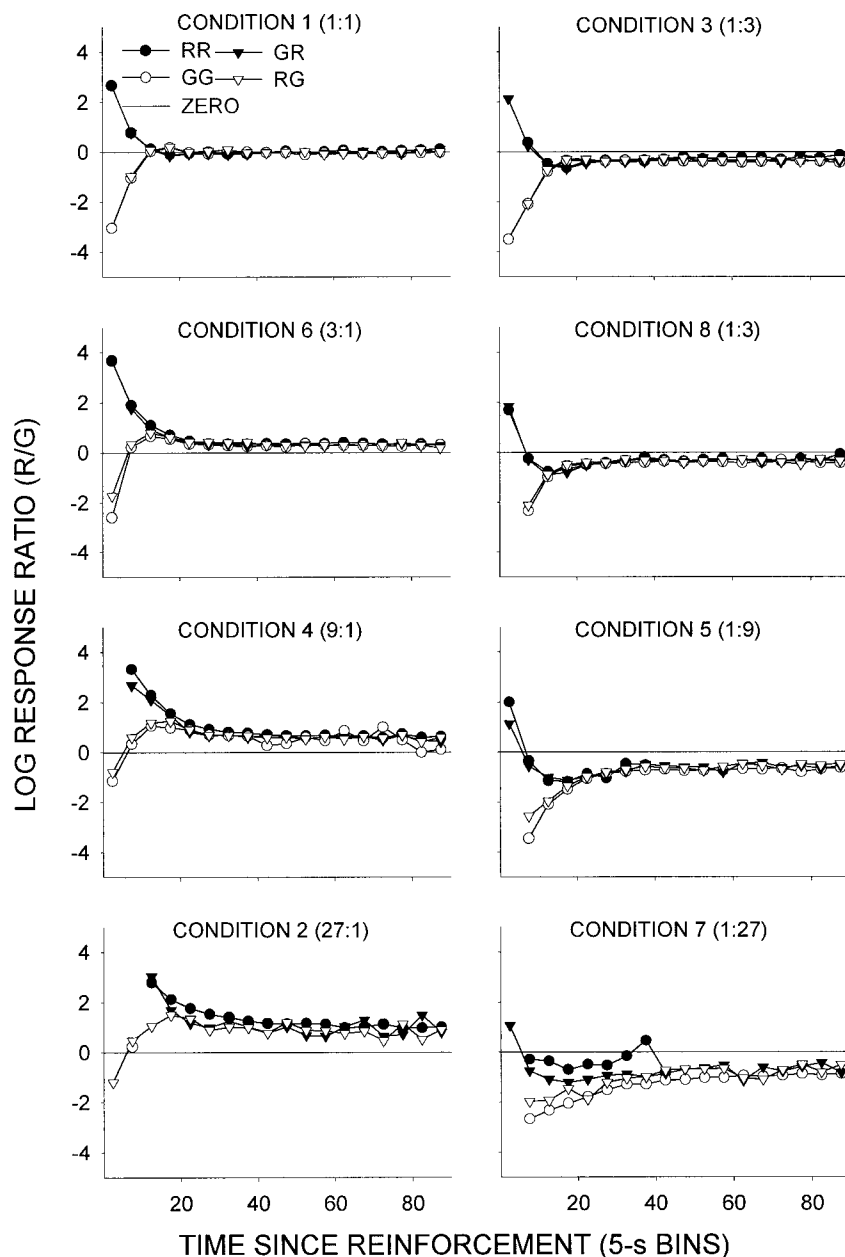


Fig. 5. The log response ratio emitted in successive 5-s time bins following the four possible two-reinforcer sequences. Also plotted are reference lines indicating zero on the y axes.

were not solely the result of rapidly changing contingencies. Many similar effects were evident in the present results. A number of analyses showed that individual reinforcers had consistent effects on behavior. Moreover, if the behavior in Condition 1 of the present study (Figure 4) is compared to the behavior in Condition 8 of the Landon and Davison

study (see their Figure 6), the effects of individual reinforcers are almost superimposable. In other conditions, reinforcer-by-reinforcer analyses showed the same large effects of disconfirmations and smaller effects of sequences of reinforcers at the same alternative (Figure 4).

The analyses of the effects of preceding re-

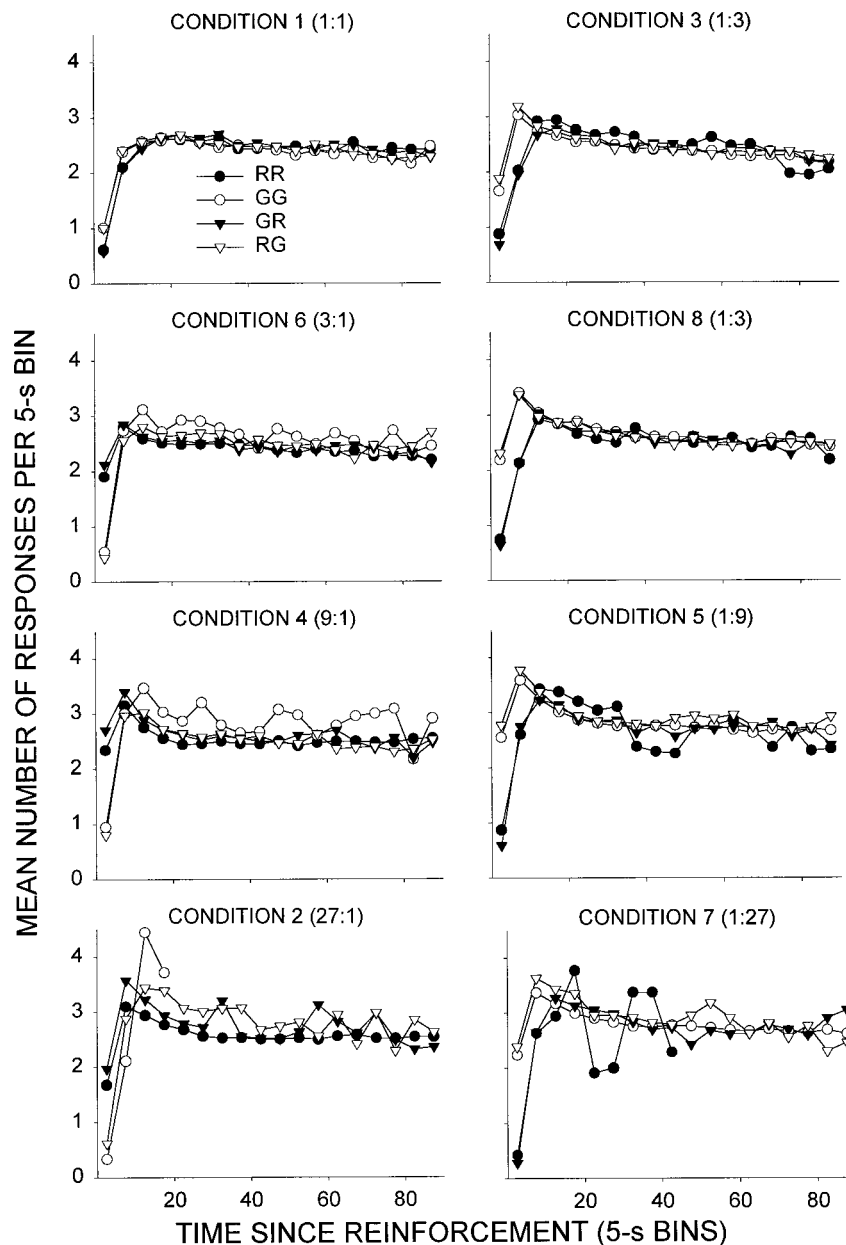


Fig. 6. The mean number of responses emitted in successive 5-s time bins following the four possible two-reinforcer sequences. The filled symbols show behavior when the most recent reinforcer was obtained at the red alternative, and the open symbols show behavior when the most recent reinforcer was obtained at the green alternative. Data were omitted when a particular time bin was reached fewer than five times.

inforcers on behavior (Figures 1 and 2) showed that each of the previous eight reinforcers had a small but generally positive effect on current preference. The only difference evident in the effects of these reinforcers was that the most recently ob-

tained reinforcer had a substantially larger effect on preference than the preceding ones. It was also the case that reinforcers obtained in conditions in which the reinforcer ratio was more extreme had larger effects (see Figure 2). This was also supported when the win-

dow size was varied (Figure 3), which showed that  $\log k$  decreased as window size was increased and that  $\log k$  decreased more quickly when the reinforcer ratio was more extreme. The latter finding suggests that control might have been more local in conditions with more extreme reinforcer ratios. The long-term effects of reinforcer ratio resemble the effect of context reported by Landon and Davison (2001). They reported an effect of range of variation, however, whereas in the present experiment the effect seems to be due to extremity of the reinforcer ratio. Whether these effects are the same is a question for future research, but it remains possible that Landon and Davison's finding was due to their arranging more or less extreme ratios in different conditions rather than ratios that varied over a wider or narrower range. This context effect was also evident when we analyzed preference during interreinforcer intervals (Figure 5). In this case, the duration of the shift in preference, rather than its magnitude, on the higher reinforcer-rate alternative increased with increasing reinforcer ratio.

The analyses shown in Figures 1 and 2 are particularly important because, in addition to showing clear short-term effects of individual reinforcers, they also provide evidence for substantially longer term effects of reinforcers. Across conditions, the constant ( $\log k$ ), which measures the effects of longer term aggregations of reinforcers, changed with the reinforcer ratio in effect. A similar effect was also evident when log response ratios were analyzed during interreinforcer intervals (Figure 5). In both these analyses, behavior stabilized after reinforcers at levels ordered in the same way as the reinforcer ratios arranged in each condition. Thus, converging evidence suggested both long- and short-term effects of reinforcers on behavior, and that control was not purely local, as suggested by Davison and Baum (2000).

The log response ratios at which preference stabilized during interreinforcer intervals in Figure 5 were less extreme than the  $\log k$  values obtained from the analysis of control by the eight most recent reinforcers (Figure 1). This occurred because the delivery of another reinforcer resets the interreinforcer interval to zero, and short intervals were much more common than long intervals.

Thus, most of the behavior from which  $\log k$  values were derived came from the left end of the functions shown in Figure 5, where preference was relatively extreme. This makes the orderly relation between the stable preferences shown in Figure 5 and reinforcer ratio more striking, because it represents control by the reinforcer ratio even when no reinforcers had been obtained for much longer than the average scheduled interreinforcer interval.

The asymmetry evident in the tree structures in the present results (Figure 4) represents a clear difference from the results reported by Davison and Baum (2000, 2002) and Landon and Davison (2001). This is probably because those authors always arranged conditions in which the average within-session reinforcer rates on the two alternatives were equal. That is, because their sets of arranged reinforcer ratios were always symmetrical, their sessional log reinforcer ratio was always zero. Here, we arranged nonzero log reinforcer ratios in most conditions, so it may not be surprising that we found asymmetrical trees (Figure 4). If this interpretation is correct, it again highlights the importance of longer term aggregations of reinforcers in addition to the short-term effects we have documented.

The trees (Figure 4) are generally consistent with the effects of the reinforcer ratios already discussed (e.g., Figure 5). In each condition, the entire tree moved towards the alternative providing the higher rate of reinforcement, with apparent limits to preference ordered in the same way as the reinforcer ratios. Figures 4 and 5, however, show apparent discrepancies in the effects of disconfirmations and successive same-alternative reinforcers. Figure 4 shows large effects of disconfirmations and comparatively small effects of successive same-alternative reinforcers at a reinforcer-by-reinforcer level. In contrast, Figure 5 shows that within interreinforcer intervals, reinforcers obtained from the richer alternative had large effects and reinforcers obtained from the leaner alternative had small effects, irrespective of whether the reinforcer was obtained from the same or the opposite alternative to the preceding reinforcer. Given this, and the fact that Davison and Baum (2000) used the term *disconfirmation* to describe a reinforcer that broke a se-

quence of same-alternative reinforcers when they arranged symmetrical distributions of reinforcers, it seems best in the present context (where the distribution of reinforcers was asymmetrical) to consider disconfirmations as reinforcers obtained at the leaner alternative.

Although Figure 4 shows choice following sequences of successive reinforcers obtained at the richer alternative, these log response ratios would again have consisted mainly of behavior from the left end of the functions shown in Figure 5, where preferences were relatively extreme and response rates were higher than those following a reinforcer at the leaner alternative (Figure 6). Because the durations of the transient movements in preference changed with the reinforcer ratio, the log response ratios during the interreinforcer intervals would also have changed, and the limits of the trees would also have changed across conditions. Following the disconfirmations shown in Figure 4, the transient changes in preference were short, and as Figure 6 shows, the response rates were also low. Thus, disconfirmations resulted in preferences similar to the stable levels shown in Figure 5.

Davison and Baum (2002), using the rapidly changing procedure outlined earlier (Davison & Baum, 2000), reported that preference moved towards indifference during periods of extinction following each component, irrespective of the reinforcer ratio in that component. This finding might, at first, seem inconsistent with the present results, which showed that the reinforcer ratio controlled behavior even after uncommonly long periods without reinforcement. Davison and Baum's procedure, however, arranged seven reinforcer ratios (1:27 through 27:1) that were symmetrical around 1:1. Because their sessional reinforcer ratio was 1:1, a shift towards indifference was also a shift towards the sessional reinforcer ratio. That is, our suggestion of some control by longer term reinforcer-ratio aggregates is also consistent with Davison and Baum's (2002) result.

Taken together, the present results and those reported by Davison and Baum (2000, 2002) and Landon and Davison (2001) suggest a need to reexamine the variables that control concurrent performance. The generalized matching law (Baum, 1974) has been

used to describe a wide variety of concurrent-schedule research (Davison & McCarthy, 1988). Given the regularities evident at this molar level of analysis, it has been assumed, often implicitly, that behavior is controlled by relatively large aggregations of reinforcers. Indeed, the present results were also described well at the molar level. The research outlined in the introduction suggested that, as reinforcer ratios were changed more rapidly, the locus of control became progressively more local; the present results suggest that this interpretation might be, to some extent, correct.

The present results also show, however, that even in steady-state concurrent-schedule performance, significant local effects of reinforcers occurred, similar to those reported for rapidly changing procedures. Landon and Davison (2001) suggested that their experiment provided some evidence that both local and global factors affect performance, and that the best way to view performance under the contingencies they arranged was in terms of dual control by both short- and long-term contingencies (see also Davis & Staddon, 1990). The present results suggest that this view also applies to steady-state performance, because both short- and long-term effects of reinforcers were evident. Moreover, the amount of control exerted by local contingencies seemed to change across conditions. Evidence was provided that these local contingencies exerted greater control when the reinforcer ratio was more extreme. The present results, therefore, suggest that neither purely short-term nor purely long-term processes can account for overall performance.

What is the appropriate level of analysis to study behavior? Both Skinner (1938) and Nevin (1984) have suggested that the appropriate level is that which produces orderly relations between an independent variable and behavior. This criterion has traditionally been met when relatively large aggregations of reinforcers and responses have been used, as was the case in the present research. The present research, however, showed that regularities were also evident both at a reinforcer-by-reinforcer level and when the data were aggregated across 5-s time bins during interreinforcer intervals. Neither local nor molar levels of analysis alone sufficed to account for all aspects of concurrent VI VI performance:



Molar analyses fail to explain the local effects of reinforcers, but the more local analyses fail to explain the longer term effects of reinforcers. Any future attempts to explain choice must consider both short- and long-term effects of reinforcers.

## REFERENCES

- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231-242.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, 32, 269-281.
- Belke, T. W., & Heyman, G. M. (1994). Increasing and signaling background reinforcement: Effect on the foreground response-reinforcer relation. *Journal of the Experimental Analysis of Behavior*, 61, 65-81.
- Davis, D. G. S., & Staddon, J. E. R. (1990). Memory for reward in probabilistic choice: Markovian and non-Markovian properties. *Behaviour*, 114, 37-63.
- Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, 74, 1-24.
- Davison, M., & Baum, W. M. (2002). Choice in a variable environment: Effects of blackout duration and extinction between components. *Journal of the Experimental Analysis of Behavior*, 77, 65-89.
- Davison, M. C., & Hunter, I. W. (1979). Concurrent schedules: Undermatching and control by previous experimental conditions. *Journal of the Experimental Analysis of Behavior*, 32, 233-244.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Ferguson, G. A. (1966). *Statistical analysis in psychology and education*. New York: McGraw-Hill.
- Findley, J. D. (1958). Preference and switching under concurrent scheduling. *Journal of the Experimental Analysis of Behavior*, 1, 123-144.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267-272.
- Hunter, I., & Davison, M. (1985). Determination of a behavioral transfer function: White noise analysis of session-to-session response-ratio dynamics on concurrent VI VI schedules. *Journal of the Experimental Analysis of Behavior*, 43, 43-59.
- Landon, J., & Davison, M. (2001). Reinforcer-ratio variation and its effects on rate of adaptation. *Journal of the Experimental Analysis of Behavior*, 75, 207-234.
- Lobb, B., & Davison, M. C. (1975). Performance in concurrent interval schedules: A systematic replication. *Journal of the Experimental Analysis of Behavior*, 24, 191-197.
- Marascuilo, L. A., & McSweeney, M. (1977). *Nonparametric and distribution-free methods for the social sciences*. Monterey, CA: Brooks/Cole.
- Nevin, J. A. (1984). Quantitative analysis. *Journal of the Experimental Analysis of Behavior*, 42, 421-434.
- Page, E. B. (1963). Ordered hypotheses for multiple treatments: A significance test for linear ranks. *Journal of the American Statistical Association*, 58, 216-230.
- Schofield, G., & Davison, M. (1997). Nonstable concurrent choice in pigeons. *Journal of the Experimental Analysis of Behavior*, 68, 219-232.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 12, 887-895.
- Taylor, R., & Davison, M. (1983). Sensitivity to reinforcement in concurrent arithmetic and exponential schedules. *Journal of the Experimental Analysis of Behavior*, 39, 191-198.
- Wearden, J. H., & Burgess, I. S. (1982). Matching since Baum (1979). *Journal of the Experimental Analysis of Behavior*, 38, 339-348.

Received July 23, 2001

Final acceptance January 12, 2002